

CLAIMS

We claim:

1. A method comprising:

determining a first region for a first line of memory to be cached in a cache, the memory having a plurality of regions including the first region over which a plurality of lines of the memory including the first line are partitioned, each of the plurality of regions having a first variable having a corresponding second variable;

comparing the first variable for any of the plurality of regions and the second variable for said first variable;

in response to comparing said first variable for any of the plurality of regions and the second variable for said first variable,

selecting as a second region one of the plurality of regions;

selecting as a second line one of the plurality of lines of the memory currently stored in the cache and partitioned to the second region;

replacing the second line currently stored in the cache with the first line; and

changing the first variable for one of the first and second regions.

2. The method of claim 1, further comprising changing the first variable for the other of said first and second regions.

3. The method of claim 1, wherein the cache comprises a plurality of cache sets, including a cache set to which the first and the second lines map and covering the first and the second regions.

4. The method of claim 1, wherein the first variable for each of the plurality of regions comprises a counter.

5. The method of claim 4, wherein changing the first variable for the second region comprises decrementing said first variable for the second region, and changing the first variable for the first region comprises incrementing said first variable for the first region.

6. The method of claim 1, wherein the second variable for the first variable for each of the plurality of regions comprises a threshold for the first variable; and

wherein comparing the first variable for any of the plurality of regions and the second variable for said first variable comprises determining whether said first variable for any of the plurality of regions is greater than said second variable for said first variable.

7. The method of claim 6, wherein selecting as the second region comprises:

determining whether the first variable for the first region is greater than the corresponding second variable for said first variable for the first region; and

in response to determining that said first variable for the first region is greater than said corresponding second variable, selecting the second region to be the first region.

8. The method of claim 6, wherein selecting as the second region comprises selecting as the second region the region for which the first variable most greatly exceeds the corresponding second variable.

9. The method of claim 6, wherein selecting as the second line one of the plurality of lines of the memory currently stored in the cache and partitioned to the second region comprises selecting as the second line a least recently used line of the plurality of lines of the memory currently stored in the cache and partitioned to the second region.

10. The method of claim 1, further comprising, otherwise,

selecting as the second line one of the plurality of lines of the memory currently stored in the cache;

selecting as the second region the region of the plurality of regions to which the second line is partitioned;

replacing the second line currently stored in the cache with the first line; and

changing the first variable for the second region.

11. The method of claim 10, further comprising changing the first variable for the first region.

12. The method of claim 10, wherein selecting as the second line one of the plurality of lines of the memory currently stored in the cache comprises selecting as the second line a least recently used line of the plurality of lines of the memory currently stored in the cache.

13. The method of claim 1, further comprising, prior to comparing the first variable for any of the plurality of regions and the second variable for said first variable,

determining whether the cache is not full;

in response to determining that the cache is not full,

adding the first line to the cache; and,

changing the first variable for the first region.

14. The method of claim 1, wherein the second variable for the first variable for each of the regions is responsive to the first variable for at least one other of the regions.

15. The method of claim 14, wherein the first variable for each of the regions is a counter, and the second variable for said first variable is responsive to the counter of at least one other of the regions.

16. The method of claim 1, wherein the second variable for the first variable for each of the regions is responsive to at least one previous value of said first variable for the first region.

17. The method of claim 16, wherein the first variable for each of the regions is a counter, and the second variable for said first variable is a previous value of the counter.

18. A system comprising:

a cache for caching lines of a plurality of lines of memory, a first line of the plurality of lines to be substituted for a second line of the plurality of lines currently stored in the cache;

a plurality of regions over which the plurality of lines are partitioned, and including a first region to which the first line is partitioned and a second region to which the second

line is partitioned;

a first variable for each region of the plurality of regions tracking a number of the plurality of lines partitioned to the region and currently stored in the cache;

a corresponding second variable for each first variable indicating a desirable maximum number of the plurality of lines partitioned to the region of the first variable to be stored in the cache; and,

a mechanism to, in response to determining that the first variable for any of the plurality of regions is greater than the corresponding second variable for the first variable, select as the second region one of the plurality of regions for which the first variable is greater than the corresponding second variable and select as the second line one of the lines currently stored in the cache and partitioned to the second region.

19. The system of claim 18, wherein the cache comprises a plurality of cache sets, including a cache set to which the first and the second lines map and covering the first and the second regions.

20. The system of claim 18, wherein the mechanism further is to, in response to determining that the first variable for any of the plurality of regions is greater than the corresponding second variable for the first variable, replace the second line currently stored in the cache with the first line.

21. The system of claim 18, wherein the mechanism further is to, in response to determining that the first variable for any of the plurality of regions is greater than the

corresponding second variable for the first variable, decrement the first variable for the second region and increment the first variable for the first region.

22. The system of claim 18, wherein the mechanism otherwise is to select as the second line one of the plurality of lines currently stored in the cache and determine as the second region the region of the plurality of regions to which the second line is partitioned.

23. The system of claim 18, wherein the plurality of regions over which the plurality of lines are partitioned are defined based on addresses of the plurality of lines.

24. The system of claim 23, further comprising a set of range registers to determine to which of the plurality of regions a line of the plurality of lines is partitioned.

25. The system of claim 18, wherein the plurality of regions over which the plurality of lines are partitioned are defined based on one or more characteristics of the plurality of lines selected from the group of characteristics essentially consisting of: data type of data stored in the plurality of lines; code type of code stored in the plurality of lines; read characteristics of the plurality of lines; and, write characteristics of the plurality of lines.

26. An article comprising:

a computer-readable medium; and,

means in the medium for substituting a first line of a plurality of lines of memory for a second line of the plurality of lines of the memory currently stored in a cache, by, in response to determining that a first variable for any of a plurality of regions over which

the plurality of lines are partitioned is greater than a corresponding second variable, selecting as a second region one of the plurality of regions for which the first variable is greater than the corresponding second variable and selecting as the second line one of the lines currently stored in the cache and partitioned to the second region.

27. The article of claim 26, wherein the cache comprises a plurality of cache sets, including a cache set to which the first and the second lines map and covering the first and the second regions.

28. The article of claim 26, wherein the means in the medium is further for, in response to determining that the first variable for any of the plurality of regions is greater than the corresponding second variable, replacing the second line currently stored in the cache with the first line.

29. The article of claim 26, wherein the medium is a recordable data storage medium.